

L Number	Hits	Search Text	DB	Time stamp
1	17240	disk adj read or read adj2 disk	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:48
2	5893	prefetch or prefetching	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:48
3	1143817	read or reading	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:48
4	1345	(prefetch or prefetching) with (read or reading)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:48
5	354	((interrupt or exception) near5 ((pre-read\$3 or preread\$3) or (pre-fetch\$3 or prefetch\$3) or (pre-load\$3 or preload\$3) or (read-ahead or readahead) or (look-ahead or lookahead)))	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:49
6	19	scatter/gather and (prefetch or prefetched or prefetching) and interrupt and (disk adj read or read adj2 disk)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:52
7	297	(disk adj read or read adj2 disk) and (prefetch or prefetching)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:52
8	131	((prefetch or prefetching) with (read or reading)) and ((disk adj read or read adj2 disk) and (prefetch or prefetching))	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:53
9	2	(((interrupt or exception) near5 ((pre-read\$3 or preread\$3) or (pre-fetch\$3 or prefetch\$3) or (pre-load\$3 or preload\$3) or (read-ahead or readahead) or (look-ahead or lookahead))) and (((prefetch or prefetching) with (read or reading)) and ((disk adj read or read adj2 disk) and (prefetch or prefetching))))	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:53
10	99	((interrupt or exception) near5 ((pre-read\$3 or preread\$3) or (pre-fetch\$3 or prefetch\$3) or (pre-load\$3 or preload\$3) or (read-ahead or readahead) or (look-ahead or lookahead))) and ((prefetch or prefetching) with (read or reading))	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:54
11	31	(disk adj read or read adj2 disk) with (prefetch or prefetching)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:54
12	0	(((interrupt or exception) near5 ((pre-read\$3 or preread\$3) or (pre-fetch\$3 or prefetch\$3) or (pre-load\$3 or preload\$3) or (read-ahead or readahead) or (look-ahead or lookahead))) and ((prefetch or prefetching) with (read or reading)) and ((disk adj read or read adj2 disk) with (prefetch or prefetching)))	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:54
13	897	711/112.ccls.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:55
14	504	711/111.ccls.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:55
15	674	711/113.ccls.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:55

16	514	711/137.ccls.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:55
17	319	711/213.ccls.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:55
18	0	((interrupt or exception) near5 ((pre-read\$3 or preread\$3) or (pre-fetch\$3 or prefetch\$3) or (pre-load\$3 or preload\$3) or (read-ahead or readahead) or (look-ahead or lookahead))) and ((prefetch or prefetching) with (read or reading))) and (((interrupt or exception) near5 ((pre-read\$3 or preread\$3) or (pre-fetch\$3 or prefetch\$3) or (pre-load\$3 or preload\$3) or (read-ahead or readahead) or (look-ahead or lookahead))) and ((prefetch or prefetching) with (read or reading))) and ((disk adj read or read adj2 disk) with (prefetch or prefetching)))	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:55
19	4	((interrupt or exception) near5 ((pre-read\$3 or preread\$3) or (pre-fetch\$3 or prefetch\$3) or (pre-load\$3 or preload\$3) or (read-ahead or readahead) or (look-ahead or lookahead))) and ((prefetch or prefetching) with (read or reading))) and 711/213.ccls.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/08/20 21:55

Terms used [prefetch](#) or [prefetching](#) and [interrupt](#) and [read ahead](#) and [disk](#) or [drive](#)

Found 4,297 of 120,398

Sort results

relevance

Save results to a Binder

Display results

expanded form

Search Tips

Open results in a new window

[Try an Advanced Search](#)

[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale 

1 A simple and efficient parallel disk mergesort

Rakesh D. Barve, Jeffrey Scott Vitter

June 1999 **Proceedings of the eleventh annual ACM symposium on Parallel algorithms and architectures**

Full text available:  [pdf\(1.32 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

2 The MIT Alewife machine: architecture and performance

Anant Agarwal, Ricardo Bianchini, David Chaiken, Kirk L. Johnson, David Kranz, John Kubiatowicz, Beng-Hong Lim, Kenneth Mackenzie, Donald Yeung

May 1995 **ACM SIGARCH Computer Architecture News, Proceedings of the 22nd annual international symposium on Computer architecture**, Volume 23 Issue 2

Full text available:  [pdf\(1.49 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Alewife is a multiprocessor architecture that supports up to 512 processing nodes connected over a scalable and cost-effective mesh network at a constant cost per node. The MIT Alewife machine, a prototype implementation of the architecture, demonstrates that a parallel system can be both scalable and programmable. Four mechanisms combine to achieve these goals: software-extended coherent shared memory provides a global, linear address space; integrated message passing allows compiler and operat ...

3 The MIT Alewife machine: architecture and performance

Anant Agarwal, Ricardo Bianchini, David Chaiken, Kirk L. Johnson, David Kranz, J. Kubiatowicz, B.-H.

Lim, K. Mackenzie, D. Yeung

August 1998 **25 years of the international symposia on Computer architecture (selected papers)**

Full text available:  [pdf\(1.58 MB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

4 Disk-directed I/O for MIMD multiprocessors

David Kotz

February 1997 **ACM Transactions on Computer Systems (TOCS)**, Volume 15 Issue 1

Full text available:  [pdf\(559.18 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Many scientific applications that run on today's multiprocessors, such as weather forecasting and seismic analysis, are bottlenecked by their file-I/O needs. Even if the multiprocessor is configured with sufficient I/O hardware, the file system software often fails to provide the available bandwidth to the application. Although libraries and enhanced file system interfaces can make a significant improvement, we believe that fundamental changes are needed in the file server software. We prop ...

Keywords: MIMD, collective I/O, disk-directed I/O, file caching, parallel I/O, parallel file system

5 File server scaling with network-attached secure disks

Garth A. Gibson, David F. Nagle, Khalil Amiri, Fay W. Chang, Eugene M. Feinberg, Howard Gobioff, Chen Lee, Berend Ozceri, Erik Riedel, David Rochberg, Jim Zelenka

June 1997 **ACM SIGMETRICS Performance Evaluation Review, Proceedings of the 1997 ACM SIGMETRICS international conference on Measurement and modeling of computer**

By providing direct data transfer between storage and client, network-attached storage devices have the potential to improve scalability for existing distributed file systems (by removing the server as a bottleneck) and bandwidth for new parallel and distributed file systems (through network striping and more efficient data paths). Together, these advantages influence a large enough fraction of the storage market to make commodity network-attached storage feasible. Realizing the technology's ful ...

6 Prefetching in segmented disk cache for multi-disk systems

Valery Soloviev

May 1996 **Proceedings of the fourth workshop on I/O in parallel and distributed systems: part of the federated computing research conference**

7 The process-flow model: examining I/O performance from the system's point of view

Gregory R. Ganger, Yale N. Patt

June 1993 **ACM SIGMETRICS Performance Evaluation Review, Proceedings of the 1993 ACM SIGMETRICS conference on Measurement and modeling of computer systems**, Volume 21 Issue 1

8 Query evaluation techniques for large databases

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

9 Input/output behavior of supercomputing applications

Ethan L. Miller, Randy H. Katz

August 1991 **Proceedings of the 1991 ACM/IEEE conference on Supercomputing**

10 Client-server computing in mobile environments

Jin Jing, Abdelsalam Sumi Helal, Ahmed Elmagarmid

June 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 2

Recent advances in wireless data networking and portable information appliances have engendered a new paradigm of computing, called mobile computing, in which users carrying portable devices have access to data and information services regardless of their physical location or movement behavior. In the meantime, research addressing information access in mobile environments has proliferated. In this survey, we provide a concrete framework and categorization of the various way ...

Keywords: application adaptation, cache invalidation, caching, client/server, data dissemination, disconnected operation, mobile applications, mobile client/server, mobile compuing, mobile data, mobility awareness, survey, system application

11 Informed prefetching and caching

R. H. Patterson, G. A. Gibson, E. Ginting, D. Stodolsky, J. Zelenka

12 RAID-II: a high-bandwidth network file server

A. L. Drapeau, K. W. Shirriff, J. H. Hartman, E. L. Miller, S. Seshan, R. H. Katz, K. Lutz, D. A. Patterson, E. K. Lee, P. M. Chen, G. A. Gibson

April 1994 **ACM SIGARCH Computer Architecture News , Proceedings of the 21ST annual international symposium on Computer architecture**, Volume 22 Issue 2

Full text available:  pdf(1.43 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In 1989, the RAID (Redundant Arrays of Inexpensive Disks) group at U. C. Berkeley built a prototype disk array called RAID-I. The bandwidth delivered to clients by RAID-I was severely limited by the memory system bandwidth of the disk array's host workstation. We designed our second prototype, RAID-H, to deliver more of the disk array bandwidth to file server clients. A custom-built crossbar memory system called the XBUS board connects the disks directly to the high-speed network, allowing data ...

13 Session summaries from the 17th symposium on operating systems principle (SOSP'99)

Jay Lepreau, Eric Eide

April 2000 **ACM SIGOPS Operating Systems Review**, Volume 34 Issue 2

Full text available:  pdf(3.15 MB)

Additional Information: [full citation](#), [index terms](#)

14 A cost-benefit scheme for high performance predictive prefetching

Vivekanand Vellanki, Ann L. Chervenak

January 1999 **Proceedings of the 1999 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available:  pdf(169.68 KB)

Additional Information: [full citation](#), [references](#), [index terms](#)

15 Informed multi-process prefetching and caching

Andrew Tomkins, R. Hugo Patterson, Garth Gibson

June 1997 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1997 ACM SIGMETRICS international conference on Measurement and modeling of computer systems**, Volume 25 Issue 1

Full text available:  pdf(2.90 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Informed prefetching and caching based on application disclosure of future I/O accesses (hints) can dramatically reduce the execution time of I/O-intensive applications. A recent study showed that, in the context of a single hinting application, prefetching and caching algorithms should adapt to the dynamic load on the disks to obtain the best performance. In this paper, we show how to incorporate adaptivity to disk load into the TIP2 system, which uses *cost-benefit analysis* to allocate g ...

16 Disk cache—miss ratio analysis and design considerations

Alan J. Smith

August 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 3

Full text available:  pdf(3.13 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The current trend of computer system technology is toward CPUs with rapidly increasing processing power and toward disk drives of rapidly increasing density, but with disk performance increasing very slowly if at all. The implication of these trends is that at some point the processing power of computer systems will be limited by the throughput of the input/output (I/O) system. A solution to this problem, which is described and evaluated in this paper, is disk cache

17 Anticipatory scheduling: a disk scheduling framework to overcome deceptive idleness in synchronous I/O

Sitaram Iyer, Peter Druschel

October 2001 **ACM SIGOPS Operating Systems Review , Proceedings of the eighteenth ACM symposium on operating systems principles**, Volume 35 Issue 5

Full text available:  pdf(1.61 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Disk schedulers in current operating systems are generally work-conserving, i.e., they schedule a request as soon as the previous request has finished. Such schedulers often require multiple outstanding requests from each process to meet system-level goals of performance and quality of service. Unfortunately, many common applications issue disk read requests in a synchronous manner, interspersing successive requests with short periods of computation. The scheduler chooses

the next request too ea ...

18 ARIMA time series modeling and forecasting for adaptive I/O prefetching

Nancy Tran, Daniel A. Reed

June 2001 **Proceedings of the 15th international conference on Supercomputing**

Full text available:  [pdf\(485.27 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



Bursty application I/O patterns, together with transfer limited storage devices, combine to create a major I/O bottleneck on parallel systems. This paper explores the use of time series models to forecast application I/O request times, then prefetching I/O requests during computation intervals to hide I/O latency. Experimental results with I/O intensive scientific codes show performance improvements compared to standard UNIX prefetching strategies.

Keywords: I/O, access pattern, modeling, prefetching, times series

19 Performance isolation: sharing and isolation in shared-memory multiprocessors

Ben Verghese, Anoop Gupta, Mendel Rosenblum

October 1998 **Proceedings of the eighth international conference on Architectural support for programming languages and operating systems**, Volume 33 , 32 Issue 11 , 5

Full text available:  [pdf\(1.69 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



Shared-memory multiprocessors (SMPs) are being extensively used as general-purpose servers. The tight coupling of multiple processors, memory, and I/O provides enormous computing power in a single system, and enables the efficient sharing of these resources. The operating systems for these machines (UNIX or Windows NT) provide very few controls for sharing the resources of the system among the active tasks or users. This unconstrained sharing model is a serious limitation for a server because the ...

20 The local disk controller

Gilbert E. Houtekamer

August 1985 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1985 ACM SIGMETRICS conference on Measurement and modeling of computer systems**,

Volume 13 Issue 2

Full text available:  [pdf\(1.02 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



The performance of the I/O subsystem in the 370-XA architecture has been improved considerably with the introduction of the new channel subsystem, as compared to the System/370 architecture. The emphasis in the 370-XA architecture is on reducing the CPU load associated with I/O, and on reducing the congestion in multi-CPU, shared systems, by redesigning the channel system. In this paper we will show that a reallocation of the control unit logic may triple the channel subsystem's ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)

 [QuickTime](#)

 [Windows Media Player](#)

 [Real Player](#)



Terms used [prefetch](#) or [prefetching](#) and [interrupt](#) and [read ahead](#) and [disk](#) or [drive](#)

Found 4,297 of 120,398

Sort results

relevance

Save results to a Binder

[Try an Advanced Search](#)

Display results

expanded form

Search Tips

[Try this search in The ACM Guide](#)

Open results in a new window

Results 21 - 40 of 200

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

21 Phoenix: a low-power fault-tolerant real-time network-attached storage device

Anindya Neogi, Ashish Raniwala, Tzi-cker Chiueh

October 1999 **Proceedings of the seventh ACM international conference on Multimedia (Part 1)**

Full text available: [pdf\(1.38 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Phoenix is a real-time network-attached storage device (NASD) that guarantees real-time data delivery to network clients even across single disk failure. The service interfaces that Phoenix provides are best-effort/real-time reads/writes based on unique object identifiers and block offsets. Data retrieval from Phoenix can be serviced in server push or client pull modes. Phoenix's real-time disk subsystem performance results from a standard cycle-based scan-order disk scheduling mechanism. H ...

22 A prefetching prototype for the parallel file systems on the Paragon

Meenakshi Arunachalam, Alok Choudhary

May 1995 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1995 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems**, Volume 23 Issue 1

Full text available: [pdf\(294.83 KB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

23 RAID: high-performance, reliable secondary storage

Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson

June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2

Full text available: [pdf\(3.60 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

24 The SPIFFI scalable video-on-demand system

Craig S. Freedman, David J. DeWitt

May 1995 **ACM SIGMOD Record , Proceedings of the 1995 ACM SIGMOD international conference on Management of data**, Volume 24 Issue 2

Full text available: [pdf\(1.43 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents a simulation study of a video-on-demand system. We present video server algorithms for real-time disk scheduling, prefetching, and buffer pool management. The performance of these algorithms is compared against the performance of simpler algorithms such as elevator and round-robin disk scheduling and global LRU buffer pool management. Finally, we show that the SPIFFI video-on-demand system scales nearly linearly as the number of disks, videos, and terminals is increased.

25

Exploiting the non-determinism and asynchrony of set iterators to reduce aggregate file I/O

latency

David C. Steere

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles**, Volume 31 Issue 5

Full text available:  pdf(1.87 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

26 UTLB: a mechanism for address translation on network interfaces

Yuqun Chen, Angelos Bilas, Stefanos N. Damianakis, Cezary Dubnicki, Kai Li

October 1998 **Proceedings of the eighth international conference on Architectural support for programming languages and operating systems**, Volume 33 , 32 Issue 11 , 5

Full text available:  pdf(1.76 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

An important aspect of a high-speed network system is the ability to transfer data directly between the network interface and application buffers. Such a *direct data path* requires the network interface to "know" the virtual-to-physical address translation of a user buffer, *i.e.*, the physical memory location of the buffer. This paper presents an efficient address translation architecture, User-managed TLB (UTLB), which eliminates system calls and device interrupts from the common co ...

27 Predictive prefetch in the Nemesis multimedia information service

H. Katseff, B. Robinson

October 1994 **Proceedings of the second ACM international conference on Multimedia**

Full text available:  pdf(793.52 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Critical to the success of future multimedia services is the ability to provide fast access to stored information via communications networks. In the Nemesis project, we focus on application control protocols for delivering stored multimedia to a user. We are exploring adaptive rate control schemes that make use of predictive prefetch of information from remote storage servers as a strategy for coping with short-term network congestion. For users with low-end equipment or slow network conne ...

28 Active disks: programming model, algorithms and evaluation

Anurag Acharya, Mustafa Uysal, Joel Saltz

October 1998 **Proceedings of the eighth international conference on Architectural support for programming languages and operating systems**, Volume 33 , 32 Issue 11 , 5

Full text available:  pdf(1.57 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Several application and technology trends indicate that it might be both profitable and feasible to move computation closer to the data that it processes. In this paper, we evaluate *Active Disk* architectures which integrate significant processing power and memory into a disk drive and allow application-specific code to be downloaded and executed on the data that is being read from (written to) disk. The key idea is to offload bulk of the processing to the diskresident processors and to us ...

29 A dual processor VAX 11/780

George H. Goble, Michael H. Marsh

April 1982 **Proceedings of the 9th annual symposium on Computer Architecture**

Full text available:  pdf(744.32 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the design of a dual processor VAX 11/780 built at the Purdue University Electrical Engineering School. It covers the conversion of a standard single processor VAX 11/780 into a dual processor system. A detailed description of hardware modifications performed and a parts list are included. The dual processor VAX is currently running a modified version of the UNIX (Fourth Berkeley Distribution) operating system. Because of licensing restrictions, operating sys ...

30 Special issue on the PAPA 2002 workshop: Disk scheduling policies with lookahead

Alexander Thomasian, Chang Liu

September 2002 **ACM SIGMETRICS Performance Evaluation Review**, Volume 30 Issue 2

Full text available:  pdf(1.08 MB)

Additional Information: [full citation](#), [abstract](#), [references](#)

Advances in magnetic recording technology have resulted in a rapid increase in disk capacities, but improvements in the mechanical characteristics of disks have been quite modest. For example the access time to random disk blocks has decreased by a mere factor of two, while disk capacities have increased by several orders of magnitude. High performance OLTP applications subject disks to a very demanding workload, since they require high access rates to randomly distributed disk blocks and gain l ...

Keywords: LOOK, SATF, SCAN, disk scheduling, scheduling policies with lookahead, simulation

31 Video Storage: System support for providing integrated services from networked multimedia storage servers

Ravi Wijayaratne, A. L. Narasimha Reddy

October 2001 **Proceedings of the ninth ACM international conference on Multimedia**

Full text available:  pdf(227.49 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper, we describe our experience in building an integrated multimedia storage system, Prism. Our current Linux-based implementation of Prism provides three levels of service: deadline guarantees for *periodic* applications, best-effort better response times for *interactive* applications and starvation-free throughput guarantees for *aperiodic* applications. Prism separates resource allocation from resource scheduling. Resource allocation is controlled across the service ...

Keywords: admission control, disk, file systems, multimedia, scheduling

32 Design of high performance RAID in real-time system

Peng Cheng, Hai Jin, Jiangling Zhang

June 1999 **ACM SIGARCH Computer Architecture News**, Volume 27 Issue 3

Full text available:  pdf(497.66 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

In this paper, we design a high performance RAID (*Redundant Arrays of Independent Disks*), called HUST_RAID, in real-time system. First, we give the architecture and the real-time process model of HUST_RAID. On the study of real-time disk I/O scheduling algorithm, we design a multiple queues real-time I/O scheduling algorithm. A new buffer management policy used in HUST_RAID is also given. It is the combination of some advanced technologies. We also introduce the design of software of the ...

Keywords: I/O scheduling algorithm, RAID, buffer management policy, parallel prefetching and flushing, real time system

33 Operating systems: Markov model prediction of I/O requests for scientific applications

James Oly, Daniel A. Reed

June 2002 **Proceedings of the 16th international conference on Supercomputing**

Full text available:  pdf(473.94 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Given the increasing performance disparity between processors and storage devices, exploiting knowledge of spatial and temporal I/O requests is critical to achieving high performance, particularly on parallel systems. Although perfect foreknowledge of I/O requests is rarely possible, even estimates of request patterns can potentially yield large performance gains. This paper evaluates Markov models to represent the spatial patterns of I/O requests in scientific codes. The paper also proposes thr ...

Keywords: I/O, Markov model, parallel computing, storage

34 Experiments with digital video playback

Richard Gerber, Ladan Gharai

May 1996 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1996 ACM SIGMETRICS international conference on Measurement and modeling of computer systems**, Volume 24 Issue 1

Full text available:  pdf(1.25 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper we describe our experiments on digital video applications, concentrating on the static and dynamic tradeoffs involved in video playback. Our results were extracted from a controlled series of 272 tests, which we ran in three stages. In the first stage of 120 tests, we used a simple player-monitor tool to evaluate the effects of various static parameters: *compression type*, *frame size*, *digitized rate*, *spatial quality* and *keyframe distribution*. The tests were carried out ...

35 The Vesta parallel file system

Peter F. Corbett, Dror G. Feitelson

August 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 3

Full text available:  pdf(649.08 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The Vesta parallel file system is designed to provide parallel file access to application programs running on multicomputers with parallel I/O subsystems. Vesta uses a new abstraction of files: a file is not a sequence of bytes, but rather it can be partitioned into multiple disjoint sequences that are accessed in parallel. The partitioning—which can also be changed dynamically—reduces the need for

synchronization and coordination during the access. Some control over the layout ...

Keywords: data partitioning, parallel computing, parallel file system

36 A survey of commercial parallel processors

Edward Gehringer, Janne Abularade, Michael H. Gulyen

September 1988 **ACM SIGARCH Computer Architecture News**, Volume 16 Issue 4

Full text available:  [pdf\(2.96 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

This paper compares eight commercial parallel processors along several dimensions. The processors include four shared-bus multiprocessors (the Encore Multimax, the Sequent Balance system, the Alliant FX series, and the ELXSI System 6400) and four network multiprocessors (the BBN Butterfly, the NCUBE, the Intel iPSC/2, and the FPS T Series). The paper contrasts the computers from the standpoint of interconnection structures, memory configurations, and interprocessor communication. Also, the share ...

37 Integrated document caching and prefetching in storage hierarchies based on Markov-chain predictions

Achim Kraiss, Gerhard Weikum

August 1998 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 7 Issue 3

Full text available:  [pdf\(603.01 KB\)](#)

Additional Information: [full citation](#), [abstract](#)

Large multimedia document archives may hold a major fraction of their data in tertiary storage libraries for cost reasons. This paper develops an integrated approach to the vertical data migration between the tertiary, secondary, and primary storage in that it reconciles speculative prefetching, to mask the high latency of the tertiary storage, with the replacement policy of the document caches at the secondary and primary storage level, and also considers the interaction of these policies with ...

Keywords: Caching, Markov chains, Performance, Prefetching, Scheduling, Stochastic modeling, Tertiary storage

38 Prediction caches for superscalar processors

James E. Bennett, Michael J. Flynn

December 1997 **Proceedings of the 30th annual ACM/IEEE international symposium on Microarchitecture**

Full text available:



[pdf\(1.02 MB\)](#)



[Publisher](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)



[Site](#)

Processor cycle times are currently much faster than memory cycle times, and this gap continues to increase. Adding a high speed cache memory allows the processor to run at full speed, as long as the data it needs is present in the cache. However, memory latency still affects performance in the case of a cache miss. Prediction caches use a history of recent cache misses to predict future misses and to reduce the overall cache miss rate. This paper describes several prediction caches, and introdu ...

Keywords: Dynamic scheduling, Memory latency, Stream buffer, Victim cache, Prediction cache

39 Improving the performance of log-structured file systems with adaptive methods

Jeanna Neefe Matthews, Drew Roselli, Adam M. Costello, Randolph Y. Wang, Thomas E. Anderson

October 1997 **ACM SIGOPS Operating Systems Review, Proceedings of the sixteenth ACM symposium on Operating systems principles**, Volume 31 Issue 5

Full text available:



[pdf\(2.18 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

40 Informed prefetching of collective input/output requests

Tara M. Madhyastha, Garth A. Gibson, Christos Faloutsos

January 1999 **Proceedings of the 1999 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available:



[pdf\(163.14 KB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

Results 21 - 40 of 200

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

This Page Blank (uspto)